arXiv:1703.00734v1 [stat.ML] 2 Mar 2017

Distributed Bayesian Matrix Factorization with Minimal Communication

Xiangju Qin‡ **Paul Blomstedt**‡ **Eemeli Leppäaho** **Pekka Parviainen** **Samuel Kaski**

Helsinki Institute for Information Technology HIIT
Department of Computer Science
Aalto University

Abstract

Bayesian matrix factorization (BMF) is a powerful tool for producing low-rank representations of matrices, and giving principled predictions of missing values. However, scaling up MCMC samplers to large matrices has proven to be difficult with parallel algorithms that require communication between MCMC iterations. On the other hand, designing communication-free algorithms is challenging due to the inherent unidentifiability of BMF solutions. We propose posterior propagation, an embarrassingly parallel inference procedure, which hierarchically introduces dependencies between data subsets and thus alleviates the unidentifiability problem.

1 Introduction

Latent factor models based on matrix factorization have in recent years become one of the most popular and successful approaches for collaborative filtering in recommender systems (Koren *et al.*, 2009). The main idea in such models is, given a matrix of observed values $\mathbf{Y} \in \mathbb{R}^{N \times D}$, to find two low-rank matrices $\mathbf{X} \in \mathbb{R}^{N \times K}$ and $\mathbf{W} \in \mathbb{R}^{D \times K}$ with $K \ll N, D$, such that

$$\mathbf{Y} \approx \mathbf{X}\mathbf{W}^\top. \tag{1}$$

In the context of recommender systems, $\mathbf{Y}$ is typically a matrix of preferences given by each of N users to a subset of the D items. The rows of the matrices $\mathbf{X}$ and $\mathbf{W}$ are then interpreted as K-dimensional user-specific and item-specific latent feature vectors, respectively. As the matrix $\mathbf{Y}$ is only partially observed, the goal of collaborative filtering is to predict unobserved user-item preferences based on observed ones. While the most widely known applications of recommender systems are in the consumer industry (e.g., recommending books, movies or music), they have also been used, for example, in the area of drug discovery to predict drug-protein bindings (e.g. Simm *et al.*, 2015).

A standard way of dealing with the high level of sparsity inherent in most real-world data sets is to model the observed values only (instead of imputing the missing values), while using regularization to avoid overfitting. A probabilistically justified regularized matrix factorization model was first introduced by Salakhutdinov and Mnih (2008b). The authors subsequently presented a fully Bayesian formulation of their model (Salakhutdinov and Mnih, 2008a), which sidesteps the difficulty of choosing appropriate values for the regularization parameters by placing a hyperprior over the hyperparameters (i.e., regularization parameters) and using Markov chain Monte Carlo (MCMC) to perform posterior inference. Additional advantages of a fully Bayesian approach include improved predictive accuracy, quantification of the uncertainty in predictions, the ability to incorporate prior knowledge, as well as flexible utilization of side-information (Adams *et al.*, 2010; Park *et al.*, 2013; Porteous *et al.*, 2010; Simm *et al.*, 2015).

Despite the appeal and many advantages of Bayesian matrix factorization (BMF), scaling up the posterior inference for industry-scale problems has proven difficult (Ahn *et al.*, 2015). Scaling up to this level calls for distributing both data and computations over many workers, and so far only very few distributed implementations of BMF have been considered in the literature. Recently, Ahn *et al.* (2015) proposed a solution based on distributed stochastic gradient Langevin dynamics. A distributed implementation using Gibbs sampling with a partitioned global address space approach has been presented by Chakroun *et al.* (2016). While a key factor in devising efficient distributed solutions is to be able to minimize communication between worker nodes, all of the present solutions require some degree of communication in between iterations.

One of the most promising directions in large-scale Bayesian computation in recent years has been *embarrassingly parallel* MCMC, a family of essentially communication-free algorithms, where the data are first

‡ These authors contributed equally to this work.

partitioned into multiple subsets and independent sampling algorithms are then run on each subset (Minsker *et al.*, 2014; Neiswanger *et al.*, 2014; Scott *et al.*, 2016; Srivastava *et al.*, 2015; Wang and Dunson, 2013; Wang *et al.*, 2014, 2015). In these algorithms, communication only takes places in a final aggregation step, which combines the subset posteriors to form an approximation to the full-data posterior.

The goal of this paper is to develop an embarrassingly parallel MCMC implementation for BMF. The main challenge with designing such algorithms for BMF models is that Equation (1) lacks a unique solution. Each subset posterior could converge to any of an infinite number of modes, making the aggregation step difficult to carry out in a meaningful way. Additionally, while typical embarrassingly parallel MCMC algorithms partition the data only with respect to rows (say), such that each subset depends on the same set of parameters, a large enough matrix factorization problem necessitates the data to be partitioned with respect to both rows and columns, effectively making different subsets dependent on a different or partially shared subsets of the parameters.

To overcome the problem of unidentifiability while minimizing the amount of communication needed, we propose *posterior propagation* (Section 3.2), a hierarchical embarrassingly parallel MCMC strategy in three phases, where the posteriors from each phase are propagated forwards and used as priors in the following phase. In the first phase, the algorithm is run for one data subset only, to obtain posteriors for the corresponding subsets of row-specific and column-specific parameters. In the second phase, the algorithm is run for all subsets which partially share parameters with the first subset. In the third phase, all remaining subsets are processed. Communication is thus limited to two instances of propagating posteriors from one phase to another, in addition to the final aggregation of subset posteriors. To enable efficient sharing of information between submodels, we propose an implementation based on Gaussian approximations (Section 4).

In addition to making predictions using the aggregated global model, we also evaluate experimentally the predictive performance of models constructed locally for each subset (Section 5). Our results suggest that the latter is often sufficient for achieving good predictive accuracy, especially when trained on more densely observed data, but appears to benefit from posterior propagation (omitting the aggregation step) to regularize the model. The general message of our paper is that low-communication distributed inference can provide faster computations with negligible effect on predictive accuracy.

2 Background

2.1 Bayesian Matrix Factorization

Let $\mathbf{Y} \in \mathbb{R}^{N \times D}$ be a partially observed data matrix and let $\mathbf{X} \in \mathbb{R}^{N \times K} = (\mathbf{x}_1, \ldots, \mathbf{x}_N)^\top$ and $\mathbf{W} = (\mathbf{w}_1, \ldots, \mathbf{w}_D)^\top \in \mathbb{R}^{D \times K}$ be matrices of unknown parameters. The general Bayesian matrix factorization (BMF) model is then specified by a likelihood

$$p(\mathbf{Y}|\mathbf{X}, \mathbf{W}) = \prod_{n=1}^{N} \prod_{d=1}^{D} \left[p\left(y_{nd} | \mathbf{x}_n^\top \mathbf{w}_d\right) \right]^{\mathbf{1}_{nd}},$$

and priors $p(\mathbf{X})$ and $p(\mathbf{W})$, with $\mathbf{1}_{nd}$ denoting an indicator function which equals 1 if the element y_{nd} is observed and 0 otherwise.

In many applications, the elements y_{nd} of the data matrix are assumed to be normally distributed, conditionally on the parameter vectors $\mathbf{x}_n$ and $\mathbf{w}_d$,

$$p\left(y_{nd} | \mathbf{x}_n^\top \mathbf{w}_d\right) = \mathcal{N}\left(y_{nd} | \mathbf{x}_n^\top \mathbf{w}_d, \tau^{-1}\right), \tag{2}$$

where τ denotes the data precision. Note that some formulations specify an individual precision τ_d for each column. Priors on the model parameters $\mathbf{X}$ and $\mathbf{W}$ are commonly specified as

$$p\left(\mathbf{X} | \mu_{\mathbf{X}}, \Lambda_{\mathbf{X}}^{-1}\right) = \prod_{n=1}^{N} \mathcal{N}_K\left(\mathbf{x}_n | \mu_{\mathbf{X}}, \Lambda_{\mathbf{X}}^{-1}\right), \tag{3a}$$

$$p\left(\mathbf{W} | \mu_{\mathbf{W}}, \Lambda_{\mathbf{W}}^{-1}\right) = \prod_{d=1}^{D} \mathcal{N}_K\left(\mathbf{w}_d | \mu_{\mathbf{W}}, \Lambda_{\mathbf{W}}^{-1}\right), \tag{3b}$$

where $\mathcal{N}_K$ denotes a K-dimensional normal distribution with covariance specified in terms of the precision matrix. The model formulation may additionally include priors on some or all of the hyperparameters $\mu_{\mathbf{X}}, \Lambda_{\mathbf{X}}, \mu_{\mathbf{W}}, \Lambda_{\mathbf{W}}$, as well as on the data precision parameters τ_d (e.g. Bhattacharya and Dunson, 2011; Salakhutdinov and Mnih, 2008a).

2.2 Embarrassingly Parallel MCMC

Consider a parametric model $p(\mathbf{Y}|\theta)$ with exchangeable observations $\mathbf{Y} = \{\mathbf{y}_1, \ldots, \mathbf{y}_N\}$ and parameter θ for which we wish to perform posterior inference using MCMC. If N is very large, the inference may be computationally too expensive to be carried out on a single machine. Embarrassingly parallel MCMC strategies aim to overcome this by partitioning the data $\mathbf{Y}$ into multiple disjoint subsets $\mathbf{Y}^{(1)} \cup \cdots \cup \mathbf{Y}^{(J)} = \mathbf{Y}$, and running independent sampling algorithms for each subset using a down-weighted prior $p(\theta)^{1/J}$. In most embarrassingly parallel MCMC algorithms, the aggregation of the obtained subset posteriors into a full-data posterior is based, in one way or another, on

the following *product density equation* (PDE),

$$p(\theta|\mathbf{Y}) \propto p(\theta)p(\mathbf{Y}|\theta) = \prod_{j=1}^{J} p(\theta)^{1/J} p\left(\mathbf{Y}^{(j)}|\theta\right). \quad (4)$$

Despite the apparent simplicity of the above equation, sampling from the PDE with satisfactory efficiency and accuracy is in general a challenging problem, since the involved densities are unknown and represented in terms of samples. For a recent overview of various subset posterior aggregation techniques, see Angelino *et al.* (2016).

3 Embarrassingly Parallel MCMC for BMF

3.1 Naive Approach

We now return to the BMF model introduced in Section 2.1. Assume that $\mathbf{Y}$ has been partitioned with respect to both rows and columns into $I \times J$ subsets $\mathbf{Y}^{(i,j)}$, $i = 1, \ldots, I$, $j = 1, \ldots, J$. The joint posterior density of the parameter matrices $\mathbf{X}$ and $\mathbf{W}$, given the partitioned data matrix $\mathbf{Y}$, then factorizes as

$$\begin{aligned} p(\mathbf{X}, \mathbf{W}|\mathbf{Y}) &\propto p(\mathbf{X})\, p(\mathbf{W})\, p(\mathbf{Y}|\mathbf{X}, \mathbf{W}) \quad (5) \\ &= \prod_{i=1}^{I} p\left(\mathbf{X}^{(i)}\right) \prod_{j=1}^{J} p\left(\mathbf{W}^{(j)}\right) \\ &\quad \times \prod_{i=1}^{I} \prod_{j=1}^{J} p\left(\mathbf{Y}^{(i,j)}|\mathbf{X}^{(i)}, \mathbf{W}^{(j)}\right). \end{aligned}$$

To develop an embarrassingly parallel algorithm for the above model, we first note that modeling each subset independently, for $i = 1, \ldots, I$ and $j = 1, \ldots, J$, yields

$$\begin{aligned} &p\left(\mathbf{X}^{(i)}, \mathbf{W}^{(j)}|\mathbf{Y}^{(i,j)}\right) \propto \quad (6) \\ &\quad p\left(\mathbf{X}^{(i)}\right) p\left(\mathbf{W}^{(j)}\right) p\left(\mathbf{Y}^{(i,j)}|\mathbf{X}^{(i)}, \mathbf{W}^{(j)}\right). \end{aligned}$$

As a first intuition, one may be tempted to replace the priors in the above equation with down-weighted variants $p\left(\mathbf{X}^{(i)}\right)^{1/I}$ and $p\left(\mathbf{W}^{(j)}\right)^{1/J}$ in order to recover the full-data posterior as a direct (naive) application of the PDE of Equation (4). However, it is well known that the parameters of latent factor models are in general unidentifiable, making the aggregation of independent subset posteriors difficult to carry out in a meaningful way. Although the unidentifiability could in principle be addressed by constraining $\mathbf{W}$ to be a lower triangular matrix (Lopes and West, 2004), the aggregation step would still remain problematic. Thus, instead of imposing constraints, we propose to handle the unidentifiability problem intrinsically and in a more general way with *posterior propagation* (PP), as described in the following subsection.

3.2 Posterior Propagation

To enforce identifiability and create dependencies between the inferences in the different subsets, we propose to use a hierarchical embarrassingly parallel MCMC strategy conducted in three successive phases, where the subset posteriors obtained in each phase are propagated forwards and used as priors in the following phase. The approach proceeds as follows (for an illustration, see Figure 1):

Inference phase 1. Inference is conducted for one subset only:

$$\begin{aligned} &p\left(\mathbf{X}^{(1)}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)}\right) \propto \quad (7) \\ &\quad p\left(\mathbf{X}^{(1)}\right) p\left(\mathbf{W}^{(1)}\right) p\left(\mathbf{Y}^{(1,1)}|\mathbf{X}^{(1)}, \mathbf{W}^{(1)}\right). \end{aligned}$$

Inference phase 2. Inference is conducted for subsets which share columns or rows with the first subset. Posterior marginals from phase 1 are used as priors for the shared parameters:

$$\begin{aligned} &p\left(\mathbf{X}^{(i)}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}\right) \propto \quad (8a) \\ &\quad p\left(\mathbf{X}^{(i)}\right) p\left(\mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)}\right) p\left(\mathbf{Y}^{(i,1)}|\mathbf{X}^{(i)}, \mathbf{W}^{(1)}\right), \end{aligned}$$

$$\begin{aligned} &p\left(\mathbf{X}^{(1)}, \mathbf{W}^{(j)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}\right) \propto \quad (8b) \\ &\quad p\left(\mathbf{X}^{(1)}|\mathbf{Y}^{(1,1)}\right) p\left(\mathbf{W}^{(j)}\right) p\left(\mathbf{Y}^{(1,j)}|\mathbf{X}^{(1)}, \mathbf{W}^{(j)}\right), \end{aligned}$$

for $i = 2, \ldots, I$ and $j = 2, \ldots, J$.

Inference phase 3. The remaining subsets are processed using posterior marginals propagated from phase 2 as priors:

$$\begin{aligned} &p\left(\mathbf{X}^{(i)}, \mathbf{W}^{(j)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}, \mathbf{Y}^{(1,j)}, \mathbf{Y}^{(i,j)}\right) \propto \quad (9) \\ &\quad p\left(\mathbf{X}^{(i)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}\right) p\left(\mathbf{W}^{(j)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}\right) \\ &\quad \times p\left(\mathbf{Y}^{(i,j)}|\mathbf{X}^{(i)}, \mathbf{W}^{(j)}\right). \end{aligned}$$

Aggregation phase. Combining all submodels and dividing away the multiply-counted propagated posterior

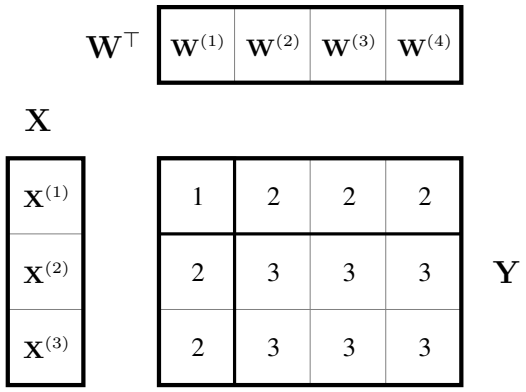

Figure 1: An example illustrating posterior propagation (PP) for a data matrix $\mathbf{Y}$ partitioned into 3×4 subsets. Subset inferences proceed in three successive phases, with posteriors obtained in one phase being propagated as priors to the next one; the numbers in the $\mathbf{Y}$ matrix denote the phase in which the particular subset is processed. Within each phase, the subsets are processed in parallel with no communication.

marginals yields the following aggregated PDE:

$$
\begin{aligned}
&p(\mathbf{X}, \mathbf{W}|\mathbf{Y}) \propto \qquad (10)\\
&\; p\left(\mathbf{X}^{(1)}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)}\right)\\
&\times \prod_{i=2}^{I} \left[p\left(\mathbf{X}^{(i)}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}\right) p\left(\mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)}\right)^{-1} \right]\\
&\times \prod_{j=2}^{J} \left[p\left(\mathbf{X}^{(1)}, \mathbf{W}^{(j)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}\right) p\left(\mathbf{X}^{(1)}|\mathbf{Y}^{(1,1)}\right)^{-1} \right]\\
&\times \prod_{i=2}^{I} \prod_{j=2}^{J} \Bigg[p\left(\mathbf{X}^{(i)}, \mathbf{W}^{(j)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}, \mathbf{Y}^{(1,j)}, \mathbf{Y}^{(i,j)}\right)\\
&\qquad \times p\left(\mathbf{X}^{(i)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}\right)^{-1}\\
&\qquad \times p\left(\mathbf{W}^{(j)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}\right)^{-1} \Bigg].
\end{aligned}
$$

To see that the above is an exact decomposition of $p(\mathbf{X}, \mathbf{W}|\mathbf{Y})$, proportional to Equation (5), it can be rewritten using the factorizations (7)–(9). While in principle, each factor of Equation (10) corresponds to a parallelly processed submodel, due to dependencies created by the propagated posteriors, the parallel inference is done in three consecutive phases.

3.3 Scalability

We will now briefly discuss the scalability of the above inference scheme in terms of total computation time relative to the number of worker nodes. Thus, with u workers available, both rows and columns can be partitioned into $\sqrt{u}+1$ parts, assuming for simplicity an equal number of partitions in both directions. This results in a total of $u + 2\sqrt{u} + 1$ subsets. Running time of a typical BMF inference algorithm per iteration is proportional to $(N + D)K^3 + MK^2$, where M is the number of observations. Thus, inference on each submodel can be done in time $t_0 := (N+D)K^3/(\sqrt{u}+1) + MK^2/(u+2\sqrt{u}+1)$, assuming equal number of observations in each subset. In phase 1, one worker does inference in time t_0. Each submodel in phase 2 can be processed in parallel, allowing completion with $2\sqrt{u}$ workers in time t_0. Finally, phase 3 is completed with u workers in time t_0. The total running time of the algorithm with u processors is therefore $3t_0 = 3[(N+D)K^3/(\sqrt{u}+1) + MK^2/(u+2\sqrt{u}+1)]$.

4 Implementation

The success of embarrassingly parallel MCMC strategies depend, in general, on the speed and accuracy of the aggregation of subposterior densities. With PP, we additionally need an efficient way of propagating information from one phase to the next. In this subsection we detail an implementation of the propagation and aggregation phases based on Gaussian approximations.

4.1 Propagation

In inference phase 1 of PP (Section 3.2), inference is done for subsets $\mathbf{X}^{(1)}$ and $\mathbf{W}^{(1)}$ of the parameters, conditional on the data subset $\mathbf{Y}^{(1,1)} \in \mathbb{R}^{N^{(1)} \times D^{(1)}}$. We first draw posterior samples from $p\left(\mathbf{X}^{(1)}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)}\right)$ using some suitable inference algorithm implemented for BMF, assuming each element of $\mathbf{Y}^{(1,1)}$ to be normally distributed (Equation (2)) and priors on $\mathbf{X}^{(1)}$ and $\mathbf{W}^{(1)}$ to be of the form given by Equations (3a) and (3b).

To enable efficient propagation of the posterior marginals $p\left(\mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)}\right)$ and $p\left(\mathbf{X}^{(1)}|\mathbf{Y}^{(1,1)}\right)$ as priors into phase 2, they are approximated as multivariate normal distributions,

$$
\widehat{p}\left(\mathbf{X}^{(1)}|\mathbf{Y}^{(1,1)}\right) = \prod_{n=1}^{N^{(1)}} \mathcal{N}_K\left(\mathbf{x}_n \mid \hat{\mu}^{(1)}_{\mathbf{x}_n}, \left[\hat{\Lambda}^{(1)}_{\mathbf{x}_n}\right]^{-1}\right),
$$

$$
\widehat{p}\left(\mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)}\right) = \prod_{d=1}^{D^{(1)}} \mathcal{N}_K\left(\mathbf{w}_d \mid \hat{\mu}^{(1)}_{\mathbf{w}_d}, \left[\hat{\Lambda}^{(1)}_{\mathbf{w}_d}\right]^{-1}\right),
$$

where $\hat{\mu}^{(1)}_{\mathbf{x}_n}, \hat{\Lambda}^{(1)}_{\mathbf{x}_n}, \hat{\mu}^{(1)}_{\mathbf{w}_d}, \hat{\Lambda}^{(1)}_{\mathbf{w}_d}$ are estimated from the posterior samples obtained for each parameter vector $\mathbf{x}_n,\ n = 1, \ldots, N^{(1)}$ and $\mathbf{w}_d,\ d = 1, \ldots, D^{(1)}$. Thus, the above priors used for each submodel of phase 2 are of the same form as in phase 1, making the same inference algorithm applicable in both phases (with the exception of inference on hyperparameters, as explained below). Finally, the posterior marginals $p\left(\mathbf{X}^{(i)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}\right)$ and

$p\left(\mathbf{W}^{(j)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}\right)$ obtained in phase 2 are approximated similarly and propagated as priors into phase 3.

For sampling of model parameters in phases 2 and 3, propagated parameters are initialized to their posterior means. Note that many implementations of the BMF model also involve sampling of the hyperparameters. While these are sampled in a standard fashion in phase 1, hyperparameters of propagated parameters will not be sampled in phases 2 and 3.

4.2 Aggregation

The aggregation of subset posteriors is carried out along marginals, making use of the approximations obtained in phases 1 and 2 of the propagation phase, which are then combined with corresponding approximations for posterior marginals from phase 3. This effectively entails an assumption of independence between the model parameters in the aggregated joint posterior, i.e. $p(\mathbf{X}, \mathbf{W}|\mathbf{Y}) = p(\mathbf{X}|\mathbf{Y})\, p(\mathbf{W}|\mathbf{Y})$, with the marginals further factorized vector-wise as

$$p(\mathbf{X}|\mathbf{Y}) = \prod_{n=1}^{N} p(\mathbf{x}_n|\mathbf{Y})$$
$$p(\mathbf{W}|\mathbf{Y}) = \prod_{d=1}^{D} p(\mathbf{w}_d|\mathbf{Y}).$$

It is worth noting, however, that the inferences carried out within each subset in phases 1–3 still result in dependent joint posteriors for the model parameters.

Assuming a multivariate normal approximation for each $p\left(\mathbf{x}_n|\mathbf{Y}\right)$, the aggregated posterior for the nth row of $\mathbf{X}$ is obtained as

$$\widehat{p}\left(\mathbf{x}_n|\mathbf{Y}\right) = \mathcal{N}_K\left(\mathbf{x}_n \mid \hat{\mu}^*_{\mathbf{x}_n}, \left[\hat{\Lambda}^*_{\mathbf{x}_n}\right]^{-1}\right), \tag{11}$$
$$\hat{\Lambda}^*_{\mathbf{x}_n} = \hat{\Lambda}^{(1)}_{\mathbf{x}_n} + \sum_{j=2}^{J}\left(\hat{\Lambda}^{(j)}_{\mathbf{x}_n} - \hat{\Lambda}^{(1)}_{\mathbf{x}_n}\right),$$
$$\hat{\mu}^*_{\mathbf{x}_n} = \left[\hat{\Lambda}^*_{\mathbf{x}_n}\right]^{-1} \times \left(\hat{\Lambda}^{(1)}_{\mathbf{x}_n}\hat{\mu}^{(1)}_{\mathbf{x}_n} + \sum_{j=2}^{J}\left(\hat{\Lambda}^{(j)}_{\mathbf{x}_n}\hat{\mu}^{(j)}_{\mathbf{x}_n} - \hat{\Lambda}^{(1)}_{\mathbf{x}_n}\hat{\mu}^{(1)}_{\mathbf{x}_n}\right)\right),$$

where $\hat{\mu}^{(j)}_{\mathbf{x}_n}, \hat{\Lambda}^{(j)}_{\mathbf{x}_n},\ j = 1, \ldots, J$, denote the estimated hyperparameters of the posterior of the jth submodel. The aggregation of each $\widehat{p}(\mathbf{w}_d|\mathbf{Y})$ is done in similar fashion.

In Equation (11), each submodel indexed by $j = 2, \ldots, J$, has the effect of the multiply counted posterior from the first submodel ($j = 1$) removed. While in principle, this will yield correctly aggregated marginals, some of the operations $\hat{\Lambda}^{(j)}_{\mathbf{x}_n} - \hat{\Lambda}^{(1)}_{\mathbf{x}_n}$ may in practice result in non-positive definite matrices, leading to numerical instabilities. In the aggregation phase, we therefore correct non-positive definite matrices by adding, to their diagonal elements, the absolute value of their smallest eigenvalue together with a small constant. The aggregation of posterior marginals using eigenvalue correction is summarized in Algorithm 1.

Algorithm 1 Posterior marginal aggregation of $p(\mathbf{x}_n|\mathbf{Y})$ using eigenvalue correction, cf. Equation (11). The function `EigenvalueCorrection`($\mathbf{M}$) finds the smallest eigenvalue of a non-positive definite symmetric matrix $\mathbf{M}$ and adds its absolute value and a small constant to all diagonal elements. The aggregation of $p(\mathbf{w}_d|\mathbf{Y})$ is done analogously.

for $j = 2 : J$ **in parallel do**
  if $\hat{\Lambda}^{(j\backslash 1)}_{\mathbf{x}_n} = \hat{\Lambda}^{(j)}_{\mathbf{x}_n} - \hat{\Lambda}^{(1)}_{\mathbf{x}_n}$ is positive definite **then**
    $\hat{\Lambda}^{*(j)}_{\mathbf{x}_n} = \hat{\Lambda}^{(j)}_{\mathbf{x}_n}$
  else
    $\hat{\Lambda}^{*(j\backslash 1)}_{\mathbf{x}_n} = \texttt{EigenvalueCorrection}\left(\hat{\Lambda}^{(j\backslash 1)}_{\mathbf{x}_n}\right)$
    $\hat{\Lambda}^{*(j)}_{\mathbf{x}_n} = \hat{\Lambda}^{*(j\backslash 1)}_{\mathbf{x}_n} + \hat{\Lambda}^{(1)}_{\mathbf{x}_n}$
  end if
end for
$\hat{\Lambda}^*_{\mathbf{x}_n} = (2-J)\hat{\Lambda}^{(1)}_{\mathbf{x}_n} + \sum_{j=2}^{J}\hat{\Lambda}^{*(j)}_{\mathbf{x}_n}$
$\hat{\mu}^*_{\mathbf{x}_n} = \left[\hat{\Lambda}^*_{\mathbf{x}_n}\right]^{-1}\left((2-J)\hat{\Lambda}^{(1)}_{\mathbf{x}_n}\hat{\mu}^{(1)}_{\mathbf{x}_n} + \sum_{j=2}^{J}\hat{\Lambda}^{*(j)}_{\mathbf{x}_n}\hat{\mu}^{(j)}_{\mathbf{x}_n}\right)$

4.3 Reordering the data

In theory, the order in which data items are processed in PP does not matter. However, for data sets which exhibit some structured pattern of sparsity (missing not at random) it may in practice be beneficial to process more densely observed subsets in the first and second phases of inference. In this way more strongly identifiable priors can be obtained for the final phase which typically has the largest number of subsets. A simple way of achieving this is to reorder the rows and columns of the data according the their proportion of observed values, which is possible to do in a distributed manner. In our experimental evaluations, we experiment with both reordered and randomly ordered data.

5 Experiments

In this section, we explore the trade-off between predictive accuracy and speed for three low-communication distributed inference schemes: (i) **posterior propagation with Gaussian approximation** (**GPP**), using the implementation of Section 4, (ii) **local independent submodels** (**LIS**), corresponding to Equation (6), and (iii) **local posterior propagation** (**LPP**), corresponding to the predictions of the local submodels of posterior propagation (without aggregation), being an intermediate between the former two.

The potential benefit of using local models for prediction (LIS and LPP) is that, while the effective number of observations used to form the parameter estimates is smaller than for the global model (GPP), the posterior distribution is not approximated using a multivariate Gaussian. Furthermore, it has all dependencies between local parameters preserved; a feature which can be of importance if the data subsets do not follow the same distribution.

5.1 Test Setup

We evaluate the above distributed inference schemes using both simulated data and two real-world data sets: MovieLens and ChEMBL. We limit our experiments to data sets that are small enough to enable a comparison with inference on the full data. In addition to inference on full data, we predict missing values using column means; this benchmark serves as a baseline and sanity check.

We evaluate the performance by predictive accuracy. To this end, we partition the data into training and test sets randomly and use root mean squared error (RMSE) on the test set as our performance measure. For prediction, we use the product of expected values of $\mathbf{X}$ and the transpose of expected values of $\mathbf{W}$. Furthermore, we use wall-clock time to measure the speedup.

For posterior inference, we use an R implementation of the Bayesian probabilistic matrix factorization (BPMF) model of Salakhutdinov and Mnih (2008a). The BPMF model considers the noise precision τ as a constant and places a normal-Wishart prior on the hyperparameters $\mu_{\mathbf{W}}$ and $\Lambda_{\mathbf{W}}$, as well as on hyperparameters $\mu_{\mathbf{X}}$ and $\Lambda_{\mathbf{X}}$. In the first phase of PP, we sample all parameters of the BPMF model. However, in the second and third phases, we sample hyperparameters $\mu_{\mathbf{W}}$ and $\Lambda_{\mathbf{W}}$ only when the posterior of $\mathbf{W}$ is not propagated. Similarly, hyperparameters $\mu_{\mathbf{X}}$ and $\Lambda_{\mathbf{X}}$ are sampled only when the posterior of $\mathbf{X}$ is not propagated. To focus on differences between parallelization schemes, we use our own implementation of BPMF in all cases and do not benchmark with other existing implementations. In principle however, any alternative implementation could be used.

In all experiments, we ran Gibbs sampling for 1200 iterations. We discarded the first 800 samples as burn-in and saved every second of the remaining samples yielding in total 200 posterior samples. The results are averaged over 5 runs.

We experiment with different partitioning schemes; a partitioning scheme $r \times c$ means that rows are partitioned into r and columns into c subsets. The partitioning scheme 1×1 refers to the full data.

We experiment with two ordering schemes. In the first scheme, we permute rows and columns randomly. In the second scheme, we reorder rows and columns into a descending order according to the proportion of observations in them. Thus, the most dense rows and columns are processed during the first two phases.

5.2 Data sets

The **MovieLens** (Harper and Konstan, 2015) data consist of movie rankings; each row corresponds to a user and each column to a movie. The data matrix has 6040 rows and 3706 columns. There are 1 000 209 observed elements. That is, about 4.5% of the elements are observed.

We follow Simm *et al.* (2015) and set $\tau = 1.5$. We analyze the performance of the models in two settings: $K = 10$ and $K = 30$. We conduct a 5-fold cross-validation study where 20% of observations are withheld and used as a test set.

The **ChEMBL** (Bento *et al.*, 2014) data describe interactions between drugs and proteins using the IC50 measure. The data set has 15 703 rows corresponding to drugs and 346 rows corresponding to proteins. The data set contains 59 280 observations which is slightly over 1% of the elements.

Again, we followed Simm *et al.* (2015) and set $\tau = 5$ and $K = 30$. Additionally, we ran experiments using $K = 10$. As ChEMBL contain only 346 columns, we did partitioning only with respect to rows.

To complement real data sets, we generated **simulated data** sets with 6040 observations and 3706 features as follows: We set the number of latent factors $K = 5$. The elements of matrices $\mathbf{W}$ and $\mathbf{X}$ were generated independently from the standard univariate normal distribution. Finally, we set $\mathbf{Y} = \mathbf{X}\mathbf{W}^\top + \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon}$ is a noise matrix whose elements were generated from a standard normal distribution. For learning, we set the parameters K and τ to the corresponding correct values, i.e. $K = 5$ and $\tau = 1$. We generated 5 independent simulated data sets.

In many real-world applications, such as collaborative filtering and the ChEMBL benchmark, the observed data are very sparse. We analyze the predictive performance of the model with respect to different levels of missing data. To this end, we randomly select 80% and 95% of the data as missing data, use these missing data as test set and the remaining data as training set. Often, the missing values are not distributed completely at random but the sparsity is structured. Thus in the second scenario, we first assign weights w_n and w_d to each row and column, respectively, such that they form an equally spaced decreasing sequence $0.9, \ldots, 0.005$. Then we assign the element y_{nd} to the test data with probability $w_n w_d$; this results in a matrix with about 80% of elements missing. This will be referred to as the structured sparsity scenario.

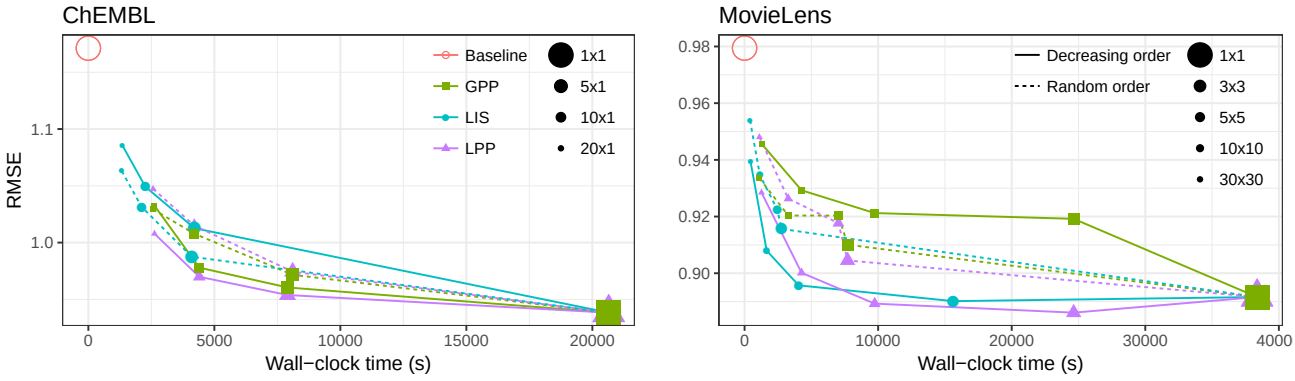

Figure 2: Test RMSE and wall-clock time on ChEMBL (left) and MovieLens data (right) with the number of latent features $K = 10$. GPP and LPP work best for ChEMBL and LIS and LPP for MovieLens.

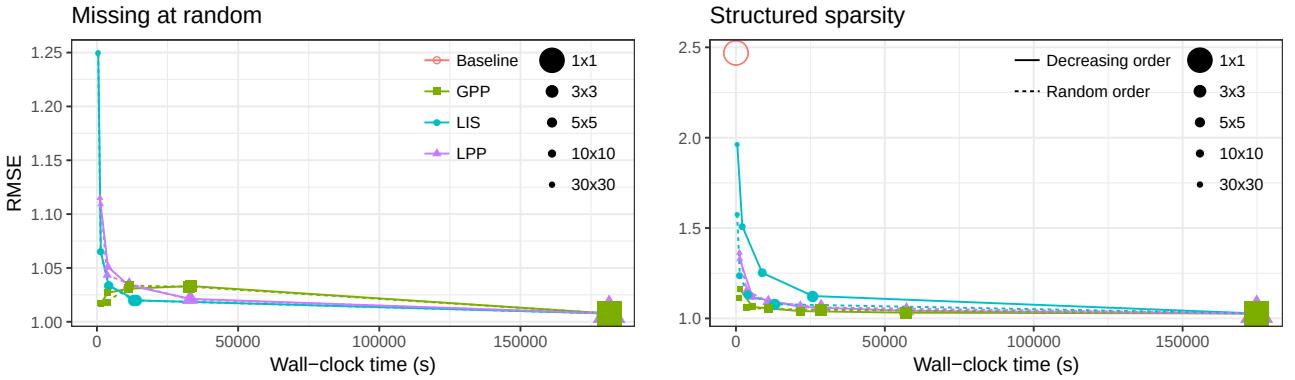

Figure 3: Test RMSE and wall-clock time on simulated data with 80% missing values. Left: values missing at random, with baseline at 2.5 RMSE omitted for clarity. Right: missing data generated with the structured sparsity scenario. Nearly horizontal curves imply that low-communication schemes speed up computation with almost no loss on accuracy.

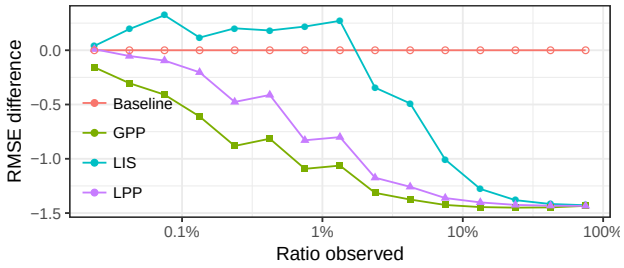

Figure 4: RMSE difference of the proposed methods to baseline on simulated data with structured sparsity, reported against the ratio of observed values in data subsets (reordered data). The posterior propagation models, especially GPP, show superior performance on sparse subsets.

5.3 Results and Conclusions

A representative set of the results for simulated data, MovieLens, and ChEMBL are shown in Figures 2 and 3; Additional figures ($K = 30$ for real data sets and 95% missing values for the simulated data) can be found in Appendix A. In the following, we summarize our conclusions and discuss the relative merits of the evaluated inference schemes.

As a general conclusion, we found that low-communication schemes can give almost an order-of-magnitude **faster computation times with a negligible effect on predictive accuracy**; this can be seen on simulated data and MovieLens (Figure 3 and right-hand side of Figure 2). Indeed, sometimes partitioning data to smaller subsets seems to even improve accuracy. Note that without approximations, PP would give the same results as the full model. The difference between them therefore tells the effect of the approximations.

We notice that GPP is more accurate than LPP on simulated data. On real data, the roles are reversed. A further analysis reveals that when we predict elements **on sparse blocks, GPP performs better** while, when predicting **on dense blocks, LPP is competitive** or sometimes even bet-

ter; see Figure 4. Intuitively, it is expected that predictions on sparse data benefit greatly on borrowing strength from other parts of the data. Thus, it is not surprising that the global model that uses more data is more accurate. However, it is surprising that the global model is sometimes very inaccurate; especially this happens with MovieLens with reordered data. However, this may be partially explained by the peculiarities of the particular data set. It seems that blockbusters behave in a different way than less popular films. It has been observed that the number of ratings per film or user is an informative feature, e.g. the winning method of the Netflix competition uses the number of rated movies per user and the number of rating users per movie as covariates (Koren, 2009).

We found that, most of the time, both **GPP and LPP are more accurate than LIS**. We also note that LIS performed best on real data. However, the test set was selected at random from the observations and thus, most of the test points lie in dense regions. It is likely that the difference between the methods would have been more pronounced in a more truthful setting where most test points are from sparse regions, cf. Figure 4. Similar behavior can be observed for the ChEMBL data set (Appendix A, Figure 7).

We observed that the **predictive accuracy of LIS is sometimes competitive**. This is not entirely surprising. Prediction with independent submodels means essentially that we infer a BMF model on a smaller data set. While using only a subset of data will lead to loss of information, if the dimension of the submodels is sufficiently large and there are sufficiently many observations then the loss in accuracy is often only moderate.

We also found that **reordering rows and columns**, in a decreasing order with respect to the number of observations, usually **improves the accuracy of GPP and LPP** compared to using a random order of rows and columns.

Recommendations. For distributed low-communication inference for BMF, we recommend posterior propagation with reordered rows and columns. Furthermore, if one wants to predict missing elements then we recommend that the choice of a scheme is made according to the density of the block where the point lies: If the block is dense then one should use LPP, otherwise one should use GPP. Note that usually the computational cost of aggregation is negligible compared to the cost of inference in submodels and thus alternating between these two schemes is essentially free.

6 Discussion

In this paper, we have introduced a hierarchical embarrassingly parallel strategy for Bayesian matrix factorization, which enables to trade off between accuracy and computational efficiency. While the present algorithm works well, it may be improved and tuned in several ways. First of all, we use a Gaussian approximation to perform propagation and aggregation. However, for aggregation there are also more sophisticated methods available (e.g. Wang *et al.*, 2015). It should be noted, though, that while more sophisticated methods typically improve accuracy, they often increase computational burden.

We have run experiments on a scale of only millions of elements. In part, this was to be able to benchmark the results of distributed inference against a model run on the full data. Additionally, scaling up to billions of observations, as required for industrial-level tasks calls for further optimization of our implementation. Note, however, that the proposed approach as such is not implementation-dependent, and it could be used together with any available well-optimized (but more communication intensive) implementation to enable further scaling with less communication.

Acknowledgements

This project has received funding from the European Unions Horizon 2020 Research and Innovation programme under Grant Agreement no. 671555 and from the Academy of Finland (grants 295503, 294238 and 292334, and Finnish Centre of Excellence in Computational Inference Research COIN). The authors gratefully acknowledge the computational resources provided by the Aalto Science-IT project.

References

Adams, R., Dahl, G., and Murray, I. (2010). Incorporating side information in probabilistic matrix factorization with Gaussian processes. In *Proceedings of the 26th Annual Conference on Uncertainty in Artificial Intelligence*, pages 1–9. AUAI Press.

Ahn, S., Korattikara, A., Liu, N., Rajan, S., and Welling, M. (2015). Large-scale distributed Bayesian matrix factorization using stochastic gradient MCMC. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 9–18.

Angelino, E., Johnson, M. J., and Adams, R. P. (2016). Patterns of scalable Bayesian inference. *arXiv preprint arXiv:1602.05221*.

Bento, A., Gaulton, A., Hersey, A., Bellis, L., Chambers, J., Davies, M., Krüger, F., Light, Y., Mak, L., McGlinchey, S., Nowotka, M., Papadatos, G., Santos, R., and Overington, J. (2014). The ChEMBL bioactivity database: an update. *Nucleic Acids Research*, **42**(D1), D1083–D1090.

Bhattacharya, A. and Dunson, D. B. (2011). Sparse Bayesian infinite factor models. *Biometrika*, **98**(2), 291–306.

Chakroun, I., Haber, T., Aa, T. V., and Kovac, T. (2016). Exploring parallel implementations of the Bayesian probabilistic matrix factorization. In *Proceedings of the 24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing*, pages 119–126.

Harper, F. M. and Konstan, J. A. (2015). The MovieLens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems*, **5**(4).

Koren, Y. (2009). The BellKor solution to the Netflix grand prize. *Netflix prize documentation*, **81**, 1–10.

Koren, Y., Bell, R., and Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, **42**(8), 30–37.

Lopes, H. F. and West, M. (2004). Bayesian model assessment in factor analysis. *Statistica Sinica*, **14**(1), 41–67.

Minsker, S., Srivastava, S., Lin, L., and Dunson, D. (2014). Robust and scalable Bayes via a median of subset posterior measures. *arXiv preprint arXiv:1403.2660*.

Neiswanger, W., Wang, C., and Xing, E. (2014). Asymptotically exact, embarrassingly parallel MCMC. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, pages 623–632. AUAI Press.

Park, S., Kim, Y.-D., and Choi, S. (2013). Hierarchical Bayesian matrix factorization with side information. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, pages 1593–1599. AAAI Press.

Porteous, I., Asuncion, A., and Welling, M. (2010). Bayesian matrix factorization with side information and Dirichlet process mixtures. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 563–568. AAAI Press.

Salakhutdinov, R. and Mnih, A. (2008a). Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. In *Proceedings of the 25th International Conference on Machine Learning*, pages 880–887. ACM.

Salakhutdinov, R. and Mnih, A. (2008b). Probabilistic matrix factorization. In *Advances in Neural Information Processing Systems 20*. MIT Press.

Scott, S. L., Blocker, A. W., Bonassi, F. V., Chipman, H. A., George, E. I., and McCulloch, R. E. (2016). Bayes and big data: the consensus Monte Carlo algorithm. *International Journal of Management Science and Engineering Management*, **11**(2), 78–88.

Simm, J., Arany, A., Zakeri, P., Haber, T., Wegner, J. K., Chupakhin, V., Ceulemans, H., and Moreau, Y. (2015). Macau: Scalable Bayesian multi-relational factorization with side information using MCMC. *arXiv preprint arXiv:1509.04610*.

Srivastava, S., Li, C., and Dunson, D. (2015). Scalable Bayes via barycenter in Wasserstein space. *arXiv preprint arXiv:1508.05880*.

Wang, X. and Dunson, D. (2013). Parallelizing MCMC via Weierstrass sampler. *arXiv preprint arXiv:1312.4605*.

Wang, X., Peng, P., and Dunson, D. B. (2014). Median selection subset aggregation for parallel inference. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2195–2203. Curran Associates, Inc.

Wang, X., Guo, F., Heller, K. A., and Dunson, D. B. (2015). Parallelizing MCMC with random partition trees. In *Advances in Neural Information Processing Systems 28*, pages 451–459. Curran Associates, Inc.

A Additional Figures

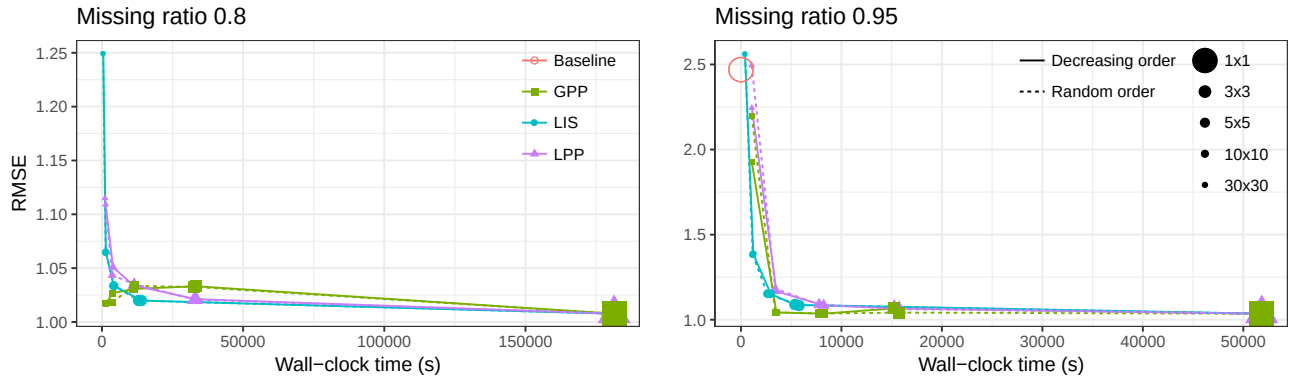

Figure 5: Test RMSE and wall-clock time of the proposed models on simulated data, reported both for data in decreasing and random order. Left: data missing at random with 80% missing data, with baseline at 2.5 RMSE omitted for clarity. Right: 95% of the elements missing at random.

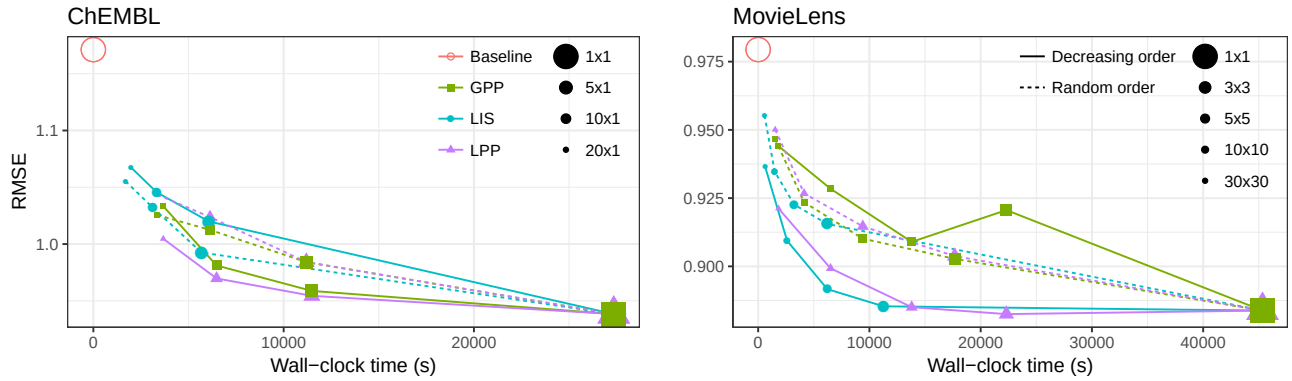

Figure 6: Test RMSE and wall-clock time on ChEMBL (left) and MovieLens data (right) with the number of latent features $K = 30$. Similar to $K = 10$, ChEMBL benefits clearly from posterior propagation with ordered data, but local independent submodels (LIS) are mostly superior on MovieLens. The best local predictions outperform the full model (1x1), implying that the ordering reveals substructure of the data set (e.g. blockbusters vs. less popular films).

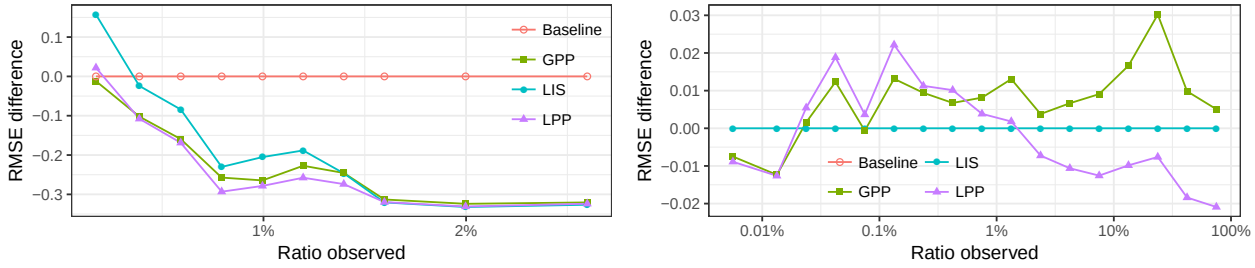

Figure 7: Relative RMSE on regions with varying density for ChEMBL (left; vs. baseline) and MovieLens (right; vs. LIS). For MovieLens the y axis is truncated for improved clarity; baseline results in around 0.2 RMSE difference across the observation ratio range. On ChEMBL data posterior propagation is shown to be advantageous for very sparsely observed subsets of data, whereas on MovieLens this is not so clear.